

AMENDMENTS TO THE CLAIMS

1 1. (Currently amended) A computer system in which a virtual memory system
2 is implemented, the computer system comprising a plurality of virtual memory pages,
3 some of which being backed by a set of one or more physical memory pages and others
4 of which being backed by a secondary memory, the computer system further comprising
5 a primary data store containing a ~~first~~ primary data item and a ~~secondary~~ derived data
6 store containing a ~~second~~ derived data item, the ~~first~~ primary data item and the ~~second~~
7 derived data item each providing memory mapping information about a first virtual
8 memory page, the ~~second~~ derived data item being derived from the ~~first~~ primary data
9 item, the computer system further comprising:

10 an actor, the actor modifying the ~~first~~ primary data item in the primary data
11 store so that the ~~second~~ derived data item in the ~~secondary~~ derived data store is
12 incoherent with the ~~first~~ primary data item;

13 a producer, the producer providing a first information indicating that the
14 ~~first~~ primary data item in the primary data store has been modified; and

15 a consumer,

16 wherein the consumer receives the first information provided by the
17 producer and, on occurrence of a coherency event at which the ~~secondary~~
18 derived data item is not to be incoherent with the primary data item, the
19 consumer ~~receives the information provided by the producer and~~ eliminates the
20 incoherency between the ~~second~~ derived data item and the ~~first~~ primary data
21 item.

1 2. (New) The computer system of claim 1, wherein the producer also provides
2 a second information indicating a derived value based on the modification that has been
3 made to the primary data item and a third information that can be used to determine
4 whether the second information remains valid, and the consumer uses the second
5 information to modify the derived data item to make the derived data item coherent with
6 the primary data item if the third information indicates that the second information is
7 valid.

1 3. (New) The computer system of claim 2, wherein, on occurrence of the
2 coherency event, if the third information indicates that the second information is no
3 longer valid, the consumer eliminates the incoherency between the derived data item
4 and the primary data item by invalidating the derived data item.

1 4. (New) The computer system of claim 2, wherein, on occurrence of the
2 coherency event, if the third information indicates that the second information is no
3 longer valid, the consumer determines a current value for the primary data item in the
4 primary data store and uses this current value for the primary data item to modify the
5 derived data item in the derived data store to make the derived data item coherent with
6 the primary data item.

1 5. (New) The computer system of claim 2, wherein the second information
2 and the third information comprise distinct data values.

1 6. (New) The computer system of claim 1, wherein the producer is the same
2 software entity as the actor.

1 7. (New) The computer system of claim 1, wherein the coherency event
2 comprises encountering a computer instruction that invalidates an entry within a
3 Translation Lookaside Buffer.

1 8. (New) The computer system of claim 1, wherein the coherency event
2 comprises a situation in which a Translation Lookaside Buffer is flushed.

1 9. (New) The computer system of claim 1, further comprising a patch channel
2 that is used by the producer to communicate the first information to the consumer.

1 10. (New) The computer system of claim 9, wherein the producer can write to
2 the patch channel at the same time that the consumer reads from the patch channel,
3 without any data races.

1 11. (New) The computer system of claim 10, wherein the patch channel is
2 stored in memory that is shared by the producer and the consumer, wherein an entry in
3 the patch channel is indicated as being in-use by the producer at a point in time that is
4 no earlier than the time when the rest of the entry is written to the patch channel, and
5 wherein the entry in the patch channel is indicated as no longer being in-use by the
6 consumer at a point in time that is no earlier than the time when the rest of the entry is
7 read from the patch channel.

1 12. (New) The computer system of claim 9, wherein the patch channel
2 comprises a half-duplex communication channel that is used exclusively by the
3 producer and the consumer.

1 13. (New) The computer system of claim 12, wherein the patch channel also
2 comprises a set data structure stored in memory that is shared by the producer and the
3 consumer.

1 14. (New) The computer system of claim 13, wherein the set data structure
2 comprises a data array containing a plurality of entries.

1 15. (New) The computer system of claim 9, further comprising additional
2 producers and additional consumers, and wherein the patch channel is one of a plurality
3 of patch channels in a patch channel matrix, wherein each patch channel in the patch
4 channel matrix is used by a different pair consisting of a single producer and a single
5 consumer.

1 16. (New) The computer system of claim 1, further comprising virtualization
2 software for supporting a virtual machine, wherein the primary data store comprises a
3 guest operating system page table containing translations from virtual memory pages to
4 guest physical pages and the derived data store comprises a shadow page table
5 containing translations from virtual memory pages to hardware physical pages.

1 17. (New) The computer system of claim 16, wherein the modification to the
2 primary data item triggers a write trace that has been placed on the primary data store,
3 and the producer provides the first information in response to the triggering of the write
4 trace.

1 18. (New) The computer system of claim 16, wherein the computer system
2 comprises a plurality of physical processors and the virtual machine comprises a
3 plurality of virtual processors, wherein a first virtual processor is the producer and a
4 second virtual processor is the consumer.

1 19. (New) The computer system of claim 16, wherein the computer system
2 comprises a plurality of physical processors and the virtual machine comprises a
3 plurality of virtual processors, wherein a first virtual processor is both the producer and
4 the consumer.

1 20. (New) A computer program embodied in a tangible medium, the computer
2 program being executable as part of a virtual computer system, the virtual computer
3 system comprising physical hardware and virtualization software for supporting a virtual
4 machine, the virtual machine running a guest operating system (OS) with a guest OS
5 page table, the guest OS page table containing address translations from virtual
6 addresses to guest physical addresses, the virtualization software generating a first
7 shadow page table derived from entries in the guest OS page table, the first shadow
8 page table containing address translations from virtual addresses to hardware physical
9 addresses, the computer program eliminating incoherencies between the address

10 translations in the guest OS page table and the address translations in the first shadow
11 page table, the guest OS page table containing a first primary entry providing an
12 address translation for a first virtual address and the first shadow page table containing
13 a first derived entry that was derived from the first primary entry in the guest OS page
14 table, the first derived entry also providing an address translation for the first virtual
15 address, the computer program comprising:

16 a first producer, the first producer responding to a first modification to the
17 first primary entry in the guest OS page table, the first modification causing an
18 incoherency between the first primary entry in the guest OS page table and the
19 first derived entry in the first shadow page table, the first producer providing a
20 first information indicating that a change has been made to a translation in the
21 guest OS page table for the first virtual address; and

22 a first consumer, the first consumer receiving the first information from the
23 first producer and, on occurrence of a coherency event at which the first shadow
24 page table is not to be incoherent with the guest OS page table with respect to
25 the first virtual address, the first consumer eliminates the incoherency between
26 the first primary entry and the first derived entry.

1 21. (New) The computer program of claim 20, wherein the first producer also
2 provides a second information indicating a derived value based on the first modification
3 that has been made to the first primary entry and a third information that can be used to
4 determine whether the second information remains valid, and the first consumer uses
5 the second information to modify the first derived entry to make the first derived entry
6 coherent with the first primary entry if the third information indicates that the second
7 information is valid.

1 22. (New) The computer program of claim 20, wherein a patch channel is
2 used exclusively by the first producer and the first consumer, the patch channel being
3 used to communicate the first information from the first producer to the first consumer,
4 and wherein an entry in the patch channel is indicated as being in-use by the first
5 producer at a point in time that is no earlier than the time when the rest of the entry is
6 written to the patch channel, and wherein the entry in the patch channel is indicated as
7 no longer being in-use by the first consumer at a point in time that is no earlier than the
8 time when the rest of the entry is read from the patch channel, so that the first producer
9 can write to the patch channel at the same time that the first consumer reads from the
10 patch channel, without any data races.

1 23. (New) The computer program of claim 20, further comprising a second
2 consumer, wherein the virtualization software also generates a second shadow page
3 table derived from entries in the guest OS page table, the second shadow page table
4 also containing address translations from virtual addresses to hardware physical
5 addresses, the computer program also eliminating incoherencies between the address
6 translations in the guest OS page table and the address translations in the second
7 shadow page table, the second shadow page table containing a second derived entry
8 that also provides an address translation for the first virtual address, wherein the second
9 derived entry was derived from the first primary entry in the guest OS page table prior to
10 the first modification of the first primary entry so that the first modification to the first
11 primary entry also causes an incoherency between the first primary entry and the
12 second derived entry, wherein the virtual computer system further comprises a patch
13 channel matrix comprising a first patch channel that is used exclusively by the first
14 producer and the first consumer and a second patch channel that is used exclusively by
15 the first producer and the second consumer, the first patch channel being used to
16 communicate the first information from the first producer to the first consumer and the
17 second patch channel being used to communicate the first information from the first
18 producer to the second consumer, the second consumer also receiving the first
19 information from the first producer and, on occurrence of a coherency event at which

20 the second shadow page table is not to be incoherent with the guest OS page table with
21 respect to the first virtual address, the second consumer eliminates the incoherency
22 between the first primary entry and the second derived entry.

1 24. (New) The computer program of claim 23, further comprising a second
2 producer, the second producer responding to a second modification to a second primary
3 entry in the guest OS page table, wherein the patch channel matrix further comprises a
4 third patch channel that is used exclusively by the second producer and the first
5 consumer and a fourth patch channel that is used exclusively by the second producer
6 and the second consumer, the third patch channel being used to communicate a fourth
7 information about the second modification from the second producer to the first
8 consumer and the fourth patch channel being used to communicate the fourth
9 information from the second producer to the second consumer.

1 25. (New) The computer program of claim 20, wherein the computer program
2 is implemented within the virtualization software.

1 26. (New) A method for eliminating an incoherency between a primary data
2 item in a primary data store and a derived data item in a derived data store, the primary
3 data store and the derived data store being in a computer system, the derived data item
4 being derived from the primary data item, the method comprising:
5 determining that a first modification is made to the primary data item in the
6 primary data store, the first modification causing the derived data item in the
7 derived data store to be incoherent with the primary data item;
8 providing a first information indicating that the primary data item in the
9 primary data store has been modified; and
10 on occurrence of a coherency event at which the derived data item is not
11 to be incoherent with the primary data item, and based on the first information,
12 eliminating the incoherency between the derived data item and the primary data
13 item.

1 27. (New) The method of claim 26, wherein the primary data item and the
2 derived data item each provide memory mapping information about a first virtual
3 memory page and the first information indicates that the primary data item has been
4 modified by specifying a physical address of the primary data item.

1 28. (New) The method of claim 26, further comprising providing a second
2 information indicating a derived value based on the first modification that has been
3 made to the primary data item and a third information that can be used to determine
4 whether the second information remains valid, and using the second information to
5 modify the derived data item to make the derived data item coherent with the primary
6 data item if the third information indicates that the second information is valid.

1 29. (New) The method of claim 26, wherein a producer provides the first
2 information to a consumer using a patch channel that is used exclusively by the
3 producer and the consumer, and the consumer eliminates the incoherency between the
4 derived data item and the primary data item.

1 30. (New) The method of claim 29, wherein the consumer retrieves the first
2 information provided by the producer and eliminates the incoherency between the
3 derived data item and the primary data item both in response to the occurrence of the
4 coherency event.

1 31. (New) The method of claim 29, wherein the incoherency between the
2 derived data item and the primary data item is eliminated by invalidating the derived
3 data item from the derived data store.

1 32. (New) The method of claim 29, wherein the incoherency between the
2 derived data item and the primary data item is eliminated by modifying the derived data
3 item to make the derived data item coherent with the primary data item.

1 33. (New) The method of claim 32, wherein, on occurrence of the coherency
2 event, the consumer determines a current value for the primary data item in the primary
3 data store and uses this current value for the primary data item to modify the derived
4 data item in the derived data store to make the derived data item coherent with the
5 primary data item.

1 34. (New) The method of claim 30, wherein the producer also provides a
2 second information indicating a derived value based on the modification that has been
3 made to the primary data item, and the consumer uses this second information to
4 modify the derived data item to make the derived data item coherent with the primary
5 data item.

1 35. (New) The method of claim 29, wherein the producer indicates that an
2 entry in the patch channel is in-use at a point in time that is no earlier than the time
3 when the rest of the entry is written to the patch channel, and wherein the consumer
4 indicates that the entry in the patch channel is no longer in-use at a point in time that is
5 no earlier than the time when the rest of the entry is read from the patch channel, so
6 that the producer can write to the patch channel at the same time that the consumer
7 reads from the patch channel, without any data races.

1 36. (New) A method for eliminating incoherencies between data items in one
2 or more primary data stores and data items in one or more derived data stores in a
3 computer system, the method comprising:
4 a first producer determining modifications that are made to data items in a
5 first primary data store, the modifications to the data items in the first primary
6 data store causing one or more incoherencies between the data items in the first
7 primary data store and data items in a first derived data store and causing one or
8 more incoherencies between the data items in the first primary data store and
9 data items in a second derived data store, the data items in the first and second

derived data stores being derived from the data items in the first primary data store;

the first producer providing information to a first consumer identifying data items in the first derived data store that are incoherent relative to the data items in the first primary data store;

the first producer providing information to a second consumer identifying data items in the second derived data store that are incoherent relative to the data items in the first primary data store;

on occurrence of a coherency event at which one or more data items in the first derived data store are not to be incoherent with corresponding data items in the first primary data store, the first consumer uses the information provided by the first producer to eliminate one or more incoherencies between the data items in the first derived data store and the data items in the first primary data store; and

on occurrence of a coherency event at which one or more data items in the second derived data store are not to be incoherent with corresponding data items in the first primary data store, the second consumer uses the information provided by the first producer to eliminate one or more incoherencies between the data items in the second derived data store and the data items in the first primary data store.

37. (New) The method of claim 36, further comprising:

a second producer determining modifications that are made to data items in a second primary data store, the modifications to the data items in the second primary data store causing one or more incoherencies between the data items in the second primary data store and data items in a third derived data store and causing one or more incoherencies between the data items in the second primary data store and data items in a fourth derived data store, the data items in the third and fourth derived data stores being derived from the data items in the second primary data store;

10 the second producer providing information to a third consumer identifying
11 data items in the third derived data store that are incoherent relative to the data
12 items in the second primary data store;

13 the second producer providing information to a fourth consumer identifying
14 data items in the fourth derived data store that are incoherent relative to the data
15 items in the second primary data store;

16 on occurrence of a coherency event at which one or more data items in
17 the third derived data store are not to be incoherent with corresponding data
18 items in the second primary data store, the third consumer uses the information
19 provided by the second producer to eliminate one or more incoherencies
20 between the data items in the third derived data store and the data items in the
21 second primary data store; and

22 on occurrence of a coherency event at which one or more data items in
23 the fourth derived data store are not to be incoherent with corresponding data
24 items in the second primary data store, the fourth consumer uses the information
25 provided by the second producer to eliminate one or more incoherencies
26 between the data items in the fourth derived data store and the data items in the
27 second primary data store.

1 38. (New) The method of claim 37, wherein the third derived data store is the
2 same as the first derived data store, the third consumer is the same as the first
3 consumer, the fourth derived data store is the same as the second derived data store,
4 the fourth consumer is the same as the second consumer, and the second primary data
5 store is the same as the first primary data store.

1 39. (New) The method of claim 38, wherein the computer system is a virtual
2 computer system comprising a virtual machine, the virtual machine comprising a first
3 virtual processor and a second virtual processor, wherein the first producer, the first
4 consumer and the third consumer perform functions related to emulating the first virtual
5 processor, and the second producer, the second consumer and the fourth consumer
6 perform functions related to emulating the second virtual processor.

1 40. (New) The method of claim 36, further comprising the first producer
2 providing additional information to the first consumer and to the second consumer
3 indicating derived values based on the modifications that have been made to the data
4 items in the first primary data store, the first consumer using the additional information
5 to modify the data items in the first derived data store to eliminate one or more
6 incoherencies between the data items in the first derived data store and the data items
7 in the first primary data store, and the second consumer using the additional information
8 to modify the data items in the second derived data store to eliminate one or more
9 incoherencies between the data items in the second derived data store and the data
10 items in the first primary data store.

1 41. (New) The method of claim 40, further comprising the first producer
2 providing second additional information to the first consumer and to the second
3 consumer that can be used to determine whether the derived values remain valid, the
4 first consumer using the second additional information to determine which derived
5 values remain valid and using the valid derived values to eliminate incoherencies
6 between the data items in the first derived data store and the data items in the first
7 primary data store, and the second consumer using the second additional information to
8 determine which derived values remain valid and using the valid derived values to
9 eliminate incoherencies between the data items in the second derived data store and
10 the data items in the first primary data store.

1 42. (New) The method of claim 36, wherein the first producer provides
2 information to the first consumer by writing the information to a first patch channel that is
3 used exclusively by the first producer and the first consumer, the first producer writing
4 the information to the first patch channel and the first consumer reading the information
5 from the first patch channel so that the first producer and the first consumer can use the
6 first patch channel at the same time, and wherein the first producer provides information
7 to the second consumer by writing the information to a second patch channel that is
8 used exclusively by the first producer and the second consumer, the first producer
9 writing the information to the second patch channel and the second consumer reading
10 the information from the second patch channel so that the first producer and the second
11 consumer can use the second patch channel at the same time.

1 43. (New) A virtual computer system, the virtual computer system comprising
2 a physical computer system with multiple physical processors, the physical computer
3 system running a virtualization software for supporting a virtual machine with multiple
4 virtual processors, the virtual machine running a guest operating system (OS) with a
5 guest OS page table, the guest OS page table containing address translations from
6 virtual addresses to guest physical addresses, the virtualization software generating a
7 first shadow page table and a second shadow page table derived from entries in the
8 guest OS page table, the first and second shadow page tables containing address
9 translations from virtual addresses to hardware physical addresses, the first shadow
10 page table being used by a first virtual processor and the second shadow page table
11 being used by a second virtual processor, the guest OS page table containing a first
12 primary entry providing an address translation for a first virtual address and the first
13 shadow page table containing a first derived entry that was derived from the first
14 primary entry in the guest OS page table, the first derived entry also providing an
15 address translation for the first virtual address, the guest OS page table also containing
16 a second primary entry providing an address translation for a second virtual address
17 and the second shadow page table containing a second derived entry that was derived
18 from the second primary entry in the guest OS page table, the second derived entry

19 also providing an address translation for the second virtual address, the virtual
20 computer system comprising:

- 21 a producer;
- 22 a first consumer;
- 23 a second consumer; and
- 24 a patch channel matrix, the patch channel matrix comprising a first patch
25 channel by which the producer communicates information to the first consumer
26 and a second patch channel by which the producer communicates information to
27 the second consumer,
- 28 wherein the producer responds to a first modification to the first primary
29 entry in the guest OS page table, the first modification causing an incoherency
30 between the first primary entry in the guest OS page table and the first derived
31 entry in the first shadow page table, the producer writing a first information to the
32 first patch channel indicating that a change has been made to a translation in the
33 guest OS page table for the first virtual address,
- 34 wherein the producer responds to a second modification to the second
35 primary entry in the guest OS page table, the second modification causing an
36 incoherency between the second primary entry in the guest OS page table and
37 the second derived entry in the second shadow page table, the producer writing
38 a second information to the second patch channel indicating that a change has
39 been made to a translation in the guest OS page table for the second virtual
40 address,
- 41 wherein the first consumer reads the first information from the first patch
42 channel and, on occurrence of a coherency event at which the first shadow page
43 table is not to be incoherent with the guest OS page table with respect to the first
44 virtual address, the first consumer eliminates the incoherency between the first
45 primary entry and the first derived entry, and
- 46 wherein the second consumer reads the second information from the
47 second patch channel and, on occurrence of a coherency event at which the
48 second shadow page table is not to be incoherent with the guest OS page table

49 with respect to the second virtual address, the second consumer eliminates the
50 incoherency between the second primary entry and the second derived entry.

1 44. (New) The virtual computer system of claim 43, wherein the producer also
2 provides to the first consumer a third information indicating a derived value based on
3 the first modification that has been made to the first primary entry and a fourth
4 information that can be used to determine whether the third information remains valid,
5 and the first consumer uses the third information to modify the first derived entry to
6 make the first derived entry coherent with the first primary entry if the fourth information
7 indicates that the third information is valid.

1 45. (New) The virtual computer system of claim 43, wherein the first patch
2 channel is used exclusively by the producer and the first consumer, and wherein an
3 entry in the patch channel is indicated as being in-use by the producer at a point in time
4 that is no earlier than the time when the rest of the entry is written to the patch channel,
5 and wherein the entry in the patch channel is indicated as no longer being in-use by the
6 first consumer at a point in time that is no earlier than the time when the rest of the entry
7 is read from the patch channel, so that the producer can write to the first patch channel
8 at the same time that the first consumer reads from the patch channel, without any data
9 races.

1 46. (New) The virtual computer system of claim 43, wherein the first virtual
2 address and the second virtual address are the same address and the first primary
3 entry and the second primary entry are the same entry.

1 47. (New) The virtual computer system of claim 43, wherein a single shadow
2 page table comprises both the first shadow page table and the second shadow page
3 table.